arXiv:1511.09389v4 [cs.DM] 26 Aug 2016

Twins in Subdivision Drawings of Hypergraphs*

René van Bevern[1,2,**], Iyad A. Kanj[3], Christian Komusiewicz[4***], Rolf Niedermeier[5], and Manuel Sorge[5]

[1] Novosibirsk State University, Novosibirsk, Russian Federation, rvb@nsu.ru
[2] Sobolev Institute of Mathematics, Siberian Branch of the Russian Academy of Sciences, Novosibirsk, Russian Federation
[3] DePaul University, Chicago, USA, ikanj@cs.depaul.edu
[4] Friedrich-Schiller-Universität Jena, Jena, Germany, christian.komusiewicz@uni-jena.de
[5] TU Berlin, Berlin, Germany, {rolf.niedermeier,manuel.sorge}@tu-berlin.de

Abstract. Visualizing hypergraphs, systems of subsets of some universe, has continuously attracted research interest in the last decades. We study a natural kind of hypergraph visualization called *subdivision drawings*. Dinkla et al. [Comput. Graph. Forum '12] claimed that only few hypergraphs have a subdivision drawing. However, this statement seems to be based on the assumption (also used in previous work) that the input hypergraph does not contain *twins*, pairs of vertices which are in precisely the same hyperedges (subsets of the universe). We show that such vertices may be necessary for a hypergraph to admit a subdivision drawing. As a counterpart, we show that the number of such "necessary twins" is upper-bounded by a function of the number m of hyperedges and a further parameter r of the desired drawing related to its number of layers. This leads to a linear-time algorithm for determining such subdivision drawings if m and r are constant; in other words, the problem is linear-time fixed-parameter tractable with respect to the parameters m and r.

1 Introduction

Hypergraph drawings are useful as visual aid in diverse applications [1], among them electronic circuit design [14] and relational databases [2, 24]. There are several methods for embedding hypergraphs in the plane. The combinatorial problem studied in this work stems from obtaining *subdivision drawings* [18, 19]. Herein, given a hypergraph $\mathcal{H}$, we divide the plane into closed regions that one-to-one correspond to the vertices of $\mathcal{H}$ in such a way that, for each hyperedge F, the union of the regions corresponding to the vertices in F is connected. Subdivision

* Appears in the *Proceedings of the 24th International Symposium on Graph Drawing and Network Visualization (GD 2016)*. Significant parts of the work were done while all authors were with TU Berlin and while RvB, IK, and MS were supported by the DFG, project NI 369/12.
** Supported by grant 16-31-60007 mol_a_dk of the Russian Foundation for Basic Research and by the Ministry of Education and Science of the Russian Federation.
*** Supported by the DFG, project KO 3669/4-1.

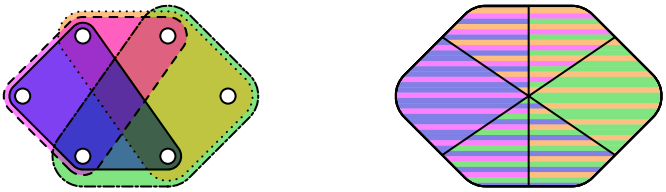

Fig. 1. Two drawings of the same hypergraph. On the left, we see a drawing in the *subset standard* in which the vertices (white circles) are enclosed by curves that correspond to hyperedges. On the right, we see a *subdivision drawing* in which we assign vertices to regions (enclosed by black lines) and we color these regions with colors that one-to-one correspond to the hyperedges; for each hyperedge, the union of the regions of the vertices in that hyperedge is connected.

drawings have also been called *vertex-based Venn diagrams* [18]. Figure 1 shows an example for such a drawing.

Subdivision drawings are a natural extension of planarity for ordinary graphs: A graph is planar if and only if it has a subdivision drawing when viewed as a hypergraph. For hypergraphs, having a subdivision drawing is a rather general concept of planar embeddings, as, for example, each Zykov planar hypergraph (meaning that the incidence graph is planar) and each hypergraph with a well-formed Euler diagram (see Flower et al. [15]) has a subdivision drawing. Still, Dinkla et al. [12] claimed that "most hypergraphs do not have [subdivision drawings]". However, this claim might have been based on the fact that several works on subdivision drawings assumed that the input hypergraph is twinless, that is, there are no two vertices contained in precisely the same hyperedges (see Mäkinen [24, p. 179], Buchin et al. [7, p. 535], and Kaufmann et al. [19, p. 399]). Twins do not seem useful at first glance: whatever role one vertex can play to obtain a subdivision drawing, its twin can also fulfill. One of our contributions is disproving the general validity of this assumption in Section 3. More specifically, we give a hypergraph with two twins that has a subdivision drawing but, removing one twin, it ceases to have one. Thus, twins may indeed be helpful to find a solution.

More generally, we can construct hypergraphs with ℓ twins that allow for subdivision drawings but cease doing so when removing one of the twins. However, the number of hyperedges in the construction grows with ℓ. It is thus natural to ask whether there is a function $\psi\colon \mathbb{N} \to \mathbb{N}$ such that, in each hypergraph with m hyperedges, we can forget all but $\psi(m)$ twins while maintaining the property of having a subdivision drawing. Using well-quasi orderings, one can relatively easily prove the *existence* of such a function ψ (see Appendix B), yet finding a closed form for ψ turned out to be surprisingly difficult: so far we could only compute a concrete upper bound when considering a second parameter r measuring the number of "layers" in the drawing. A small number r of layers, however, is a relevant special case [5, 7].

We study subdivision drawings from a combinatorial point of view, exploiting the fact that it is equivalent for a hypergraph to have a subdivision drawing and to have a support that is planar [18]. Herein, a *support* for a hypergraph $\mathcal{H} = (V, \mathcal{E})$

is a graph G on the same vertex set as $\mathcal{H}$ such that each hyperedge $F \in \mathcal{E}$ induces a connected subgraph $G[F]$. The outerplanarity number r of the support roughly translates to the number of layers in a corresponding drawing:[6] An r-outerplanar graph admits a planar embedding (without edge crossings) which has the property that, after removing r times all vertices on the outer face, we obtain an empty graph. Similar restrictions were studied before [5, 7]. Formally, we study the following problem.

Problem (r-Outerplanar Support).
Input: A connected hypergraph $\mathcal{H}$ with n vertices and m hyperedges, and $r \in \mathbb{N}$.
Question: Does $\mathcal{H}$ admit an r-outerplanar support?

Our main result is a concrete upper bound on the number $\psi(m, r)$ of twins that might be necessary to obtain an r-outerplanar support. Since superfluous twins can then be removed in linear time, this gives the following algorithmic result.

Theorem 1. *There is an algorithm solving* r-Outerplanar Support *which, for constant* r *and* m*, has linear running time.*

In contrast to Theorem 1, r-Outerplanar Support remains NP-complete for $r = \infty$ [18] and even for every fixed $r > 1$ [7] (see below). The constants in the running time of the algorithm in Theorem 1 have a large dependence on m and r. However, it is conceivable that the parameters m and r are small in practical instances: for a large number m of hyperedges, it is plausible that we obtain only hardly legible drawings unless the hyperedges adhere to some special structure. Thus, it makes sense to design algorithms particularly for hypergraphs with a small number of hyperedges, as done by Verroust and Viaud [28]. Moreover, a small outerplanarity number r leads to few layers in the drawing which may lead to aesthetically pleasing drawings, similarly to path- or cycle-supports [7].

Related work. For specifics on the relations of some different planar embeddings for hypergraphs, see Kaufmann et al. [19], Brandes et al. [5].

As mentioned before, Johnson and Pollak [18] showed that finding a *planar* support is NP-complete. Buchin et al. [7] proved that r-Outerplanar Support is NP-complete for $r = 2, 3$. From their proof it follows that r-Outerplanar Support is also NP-complete for every $r > 3$. This is due to a property of the reduction that Buchin et al. use. Given a formula ϕ in 3CNF, they construct a hypergraph $\mathcal{H}$ that has a *planar* support if and only if ϕ is satisfiable. Due to the way in which $\mathcal{H}$ is constructed, if there is any planar support, then it is 3-outerplanar. Thus, deciding whether an r-outerplanar support, $r \geq 3$, exists also decides the satisfiability of the corresponding formula.

Towards determining the computational complexity of finding an *outerplanar* hypergraph support, Brandes et al. [5] gave a polynomial-time algorithm for cactus supports (graphs in which each edge is contained in at most one cycle).

[6] We refer to Kaufmann et al. [19] for a method to obtain a subdivision drawing from a planar support.

They also showed that finding an outerplanar support (or planar support) can be done in polynomial time if, in the input hypergraph, each intersection or difference of two hyperedges is either a singleton or again a hyperedge in the hypergraph. Getting even more special, a tree support can be found in linear time [2, 26]. Buchin et al. [7] gave a polynomial-time algorithm that can deal with an additional upper bound on the vertex degrees in the tree support. Klemz et al. [20] studied so-called area-proportional Euler diagrams, for which the corresponding computational problem reduces to finding a minimum-weight tree support. Such supports can also be found in polynomial time [20, 21].

In a wider scope, motivated by drawing metro maps and metro map-like diagrams, Brandes et al. [6] studied the problem of finding *path-based* planar hypergraph supports, that is, planar supports that fulfill the additional constraint that the subgraph induced by each hyperedge contains a Hamiltonian path, giving NP-hardness and tractability results. Finding path-based tree supports is also known as the Graph Realization problem, for which several polynomial-time algorithms were already known [4].

Chen et al. [8] showed that for obtaining minimum-edge supports (not necessarily planar), twins show a similar behavior as for r-outerplanar supports: Removing a twin can increase the minimum number of edges needed for a support and finding a minimum-edge support is linear-time solvable for a constant number of hyperedges via removing superfluous twins. The proof is quite different, however.

Organization. In Section 2 we provide some technical preliminaries used throughout the work. In Section 3 we give an example that shows that twins can be crucial for a hypergraph to have a planar support. As mentioned, for each $m \in \mathbb{N}$, there is a number $\psi(m)$ such that in each hypergraph with a planar support we can safely forget all but $\psi(m)$ twins (see Appendix B). In Section 4 we give a concrete upper bound for $\psi(m)$ in the case of r-outerplanar supports and derive the linear-time algorithm for r-Outerplanar Support promised in Theorem 1. We conclude and give some directions for future research in Section 5.

2 Preliminaries

General notation. By $A \uplus B$ we denote the union of two disjoint sets A and B. For a family of sets $\mathcal{F}$, we write $\bigcup \mathcal{F}$ in place of $\bigcup_{S \in \mathcal{F}} S$. For equivalence relations ρ over some set S and $v \in S$ we use $[v]_\rho$ to denote the equivalence class of v in ρ.

Hypergraphs. A *hypergraph* $\mathcal{H}$ is a tuple $(V, \mathcal{E})$ consisting of a *vertex set* V, also denoted $V(\mathcal{H})$, and a *hyperedge set* $\mathcal{E}$, also denoted $\mathcal{E}(\mathcal{H})$. The hyperedge set $\mathcal{E}$ is a family of subsets of V, that is, $F \subseteq V$ for every hyperedge $F \in \mathcal{E}$. Where it is not ambiguous, we use $n := |V|$ and $m := |\mathcal{E}|$. When specifying running times, we use $|\mathcal{H}|$ to denote $|V(\mathcal{H})| + \sum_{F \in \mathcal{E}(\mathcal{H})} |F|$. The *size* $|F|$ of a hyperedge F is the number of vertices in it. Unless stated otherwise, we assume that hypergraphs do not contain hyperedges of size at most one or multiple copies of the same hyperedge. (These do not play any role for the problem under consideration, and removing them can be done easily and efficiently.)

A vertex $v \in V$ and a hyperedge $F \in \mathcal{E}$ are *incident* with one another if $v \in F$. For a vertex $v \in V(\mathcal{H})$, we let $\mathcal{E}_{\mathcal{H}}(v) \coloneqq \{F \in \mathcal{H} \mid v \in F\}$. If it is not ambiguous, then we omit the subscript $\mathcal{H}$ from $\mathcal{E}_{\mathcal{H}}$. A vertex v *covers* a vertex u if $\mathcal{E}(u) \subseteq \mathcal{E}(v)$. Two vertices $u, v \in V$ are *twins* if $\mathcal{E}(v) = \mathcal{E}(u)$. Clearly, the relation τ on V defined by $\forall u, v \in V\colon (u, v) \in \tau \iff \mathcal{E}(u) = \mathcal{E}(v)$ is an equivalence relation. The equivalence classes $[u]_\tau$, $u \in V$, are called *twin classes*.

Removing a vertex subset $S \subseteq V(\mathcal{H})$ from a hypergraph $\mathcal{H} = (V, \mathcal{E})$ results in the hypergraph $\mathcal{H} - S \coloneqq (V \setminus S, \mathcal{E}')$ where $\mathcal{E}'$ is obtained from $\{F \setminus S \mid F \in \mathcal{E}\}$ by removing empty and singleton sets. For brevity, we also write $\mathcal{H} - v$ instead of $\mathcal{H} - \{v\}$. The *subhypergraph shrunken to* $V' \subseteq V$ is the hypergraph $\mathcal{H}|_{V'} \coloneqq \mathcal{H} - (V \setminus V')$.

Graphs. Our notation related to graphs is basically standard and heavily borrows from Diestel's book [11]. In particular, a *bridge* of a graph is an edge whose removal increases its number of connected components. Analogously, a *cut-vertex* is a vertex whose removal increases its number of connected components. Some special notation including the gluing of graphs is given below. We use the usual notation for planar and plane graphs. An r-outerplanar graph admits a planar embedding which has the property that, after r times of removing all vertices on the outer face, we obtain an empty graph. The ith layer L_i of a plane graph is defined as the set of vertices on the outer face, after having $i - 1$ times removed all vertices on the outer face.

Boundaried graphs and gluing. For a nonnegative integer $b \in \mathbb{N}$, a *b-boundaried graph* is a tuple (G, B, β) where G is a graph, $B \subseteq V(G)$ such that $|B| = b$, and $\beta\colon B \to \{1, \dots, b\}$ is a bijection. Vertex subset B is called the *boundary* and β the *boundary labeling*. For ease of notation we also refer to (G, B, β) as the *b-boundaried* graph G with boundary B and boundary labeling β. For brevity, we also denote by *β-boundaried graph* G that b-boundaried graph G whose boundary is the domain of β and whose boundary labeling is β.

For a nonnegative integer b, the *gluing* operation $\circ_b$ maps two b-boundaried graphs to an ordinary graph as follows: Given two b-boundaried graphs G_1, G_2 with corresponding boundaries B_1, B_2 and boundary labelings β_1, β_2, to obtain the graph $G_1 \circ_b G_2$ take the disjoint union of G_1 and G_2, and identify each $v \in B_1$ with $\beta_2^{-1}(\beta_1(v)) \in B_2$. We omit the index b in $\circ_b$ if it is clear from the context.

3 Beware of removing twins

In Figure 2, we provide a concrete example that shows that twins can be necessary to obtain a 2-outer-planar support:

The vertex set of the hypergraph $\mathcal{H}$ shown in Figure 2 is $V \coloneqq \{a, b, c, d, v_a, v_b, v_d, u_b, u_c, u_d, t, t'\}$. We choose the hyperedges in such a way that t and t' are twins and $\mathcal{H}$ has a planar (more precisely, 2-outerplanar) support but $\mathcal{H} - t$ does not. First, we add to the set of hyperedges $\mathcal{E}$ of $\mathcal{H}$ the size-two hyperedges represented by solid lines between the corresponding vertices in Figure 2. The

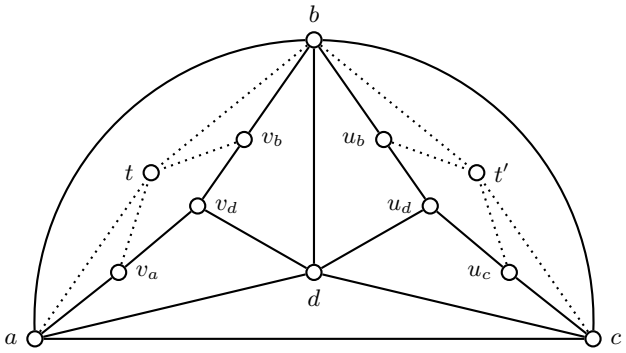

Fig. 2. A hypergraph $\mathcal{H}$ and its support, showing that twins can be essential for obtaining a 2-outer-planar support. The set of hyperedges consists of size-two hyperedges that are drawn as solid lines between the corresponding vertices and, additionally, $\{a, v_a, t, t', c\}$, $\{a, v_b, t, t', c\}$, $\{b, v_a, t, t', c\}$, $\{b, v_b, t, t', c\}$, $\{b, u_b, t, t', a\}$, $\{b, u_c, t, t', a\}$, $\{c, u_b, t, t', a\}$, and $\{c, u_c, t, t', a\}$. Note that the vertices t and t' are twins. The hypergraph $\mathcal{H}$ has a (2-outer)planar support whose edges are indicated by the solid and the dotted lines. However, $\mathcal{H} - t$ does not have a planar support.

corresponding "solid" hyperedges incident with (and only with) a, b, c, d form a K_4 and have the purpose of essentially fixing the embedding of each support G: Since the complete graph on four vertices, K_4, is 3-connected, it has only one planar embedding up to the choice of the outer face [27, p. 747]. The remaining solid hyperedges (incident with v_a, v_b, v_d and u_a, u_c, u_d) have the purpose of anchoring the u- and v-vertices within two different faces of the embedding of the K_4: These hyperedges form two connected components that are adjacent to a, b, d and b, c, d, respectively. Hence, these connected components reside in those (unique) faces of the K_4 that are incident with a, b, d and b, c, d, respectively.

With the following additional hyperedges, our goal is to enforce that t and t' are used as conduits to connect the v-vertices to c via both a and b, and to connect the u-vertices to a via both b and c. As we explain below, this is achieved by the following hyperedges:

$$\{a, v_a, t, t', c\}, \quad \{a, v_b, t, t', c\}, \quad \{b, v_a, t, t', c\}, \quad \{b, v_b, t, t', c\},$$
$$\{b, u_b, t, t', a\}, \quad \{b, u_c, t, t', a\}, \quad \{c, u_b, t, t', a\}, \quad \{c, u_c, t, t', a\}.$$

Clearly, t and t' are twins. As can easily be verified, adding t and t' and the dotted edges in Figure 2 to the graph induced by the solid edges gives a planar support for $\mathcal{H}$.

We now show that t and t' have to reside in different faces for each planar support G for $\mathcal{H}$. First, observe that, in G, either v_a is not adjacent to b or v_b is not adjacent to a. Moreover, neither of v_a and v_b is adjacent to c. Thus, to connect the subgraphs induced by the hyperedges that contain v_a or v_b, either vertex t or its twin t' must be adjacent to one of the two vertices in G. For the same reason, one of t and t' must be adjacent to one of u_b and u_c. Since there is no face in G that is simultaneously incident with one of v_a or v_b and one of u_b

or u_c, the twins t and t' thus have to be in different faces. This implies that it is impossible to obtain a planar support if t or t' is missing. Consequently, removing one vertex of a twin class can transform a yes-instance of r-Outerplanar Support into a no-instance.

To show that the example above is not a pathology of having only one pair of twins, in Appendix A we extend it so that an arbitrarily large set of twins is required for the existence of a planar support.

4 Relevant twins for r-outerplanar supports

In this section, we show that there is an explicit function ψ such that, out of each twin class of a given hypergraph $\mathcal{H}$, we can remove all but $\psi(m, r)$ twins such that the resulting hypergraph has an r-outerplanar support if and only if $\mathcal{H}$ has. In other words, we prove that the following data reduction rule is correct.

Rule 1. Let $\mathcal{H}$ be a hypergraph with m edges. If there is a twin class with more than $\psi(m, r) = 2^{6r \cdot 2^{m \cdot (2r^2+r+1)} \cdot (r+1)^{32r^2+8r}}$ vertices, then remove one vertex out of this class.

Assuming that Rule 1 is correct, Theorem 1 follows.

Proof (Theorem 1). Rule 1 can be applied exhaustively in linear time because the twin classes can be computed in linear time [17]. After this, each twin class contains at most $\psi(m, r)$ vertices, meaning that, overall, at most $2^m \psi(m, r)$ vertices remain. Testing all possible planar graphs for whether they are a support for the resulting hypergraph thus takes constant time if m and r are constant. Hence, the overall running time is linear in the input size. □

We mention in passing that, in the terms of parameterized algorithmics, exhaustive application of Rule 1 can be seen as a problem kernel (see [9], for example).

The correctness proof for Rule 1 consists of two parts. First, in Theorem 2, we show that each r-outerplanar graph has a long sequence of nested separators. Here, *nested* means that each separator separates the graph into a *left* side and a *right* side, and each left side contains all previous left sides. Furthermore, the sequence of separators has the additional property that, for any pair of separators S_1, S_2, we can glue the left side of S_1 and the right side of S_2, obtaining another r-outerplanar graph.

In the second part of the proof, we fix an initial support for our input hypergraph. We then show that, in a long sequence of nested separators for this support as above, there are two separators such that we can carry out the following procedure. We discard all vertices between the separators, glue their left and right sides, and reattach the vertices which we discarded as degree-one vertices. Furthermore, we can do this in such a way that the resulting graph is an r-outerplanar support. The reattached degree-one vertices hence are not crucial to obtain an r-outerplanar support. We will show that if our input hypergraph is large enough, that is, larger than some function of m and r, then there is always at least one non-crucial vertex which can be removed.

We now formalize our approach. Theorem 2 will guarantee the existence of a long sequence of gluable separators; it is proven in Appendix D. To formally state it, we need the following notation.

Definition 1. *For an edge bipartition $A \uplus B = E(G)$ of a graph G, let $M(A,B)$ be the set of vertices in G which are incident with both an edge in A and in B, that is,*

$$M(A,B) \;:=\; \{v \in V(G) \mid \exists a \in A \exists b \in B \colon v \in a \cap b\}.$$

We call $M(A,B)$ the middle set *of A, B. For an edge set $A \subseteq E(G)$, denote by $G\langle A\rangle := (\bigcup_{e\in A} e, A)$ the subgraph induced by A.*

Recall from Section 2 the definitions of graph gluing, boundary, and boundary labeling.

Theorem 2 ($\star$[7]). *For every connected, bridgeless, r-outerplanar graph G with n vertices there is a sequence $((A_i, B_i, \beta_i))_{i=1}^{s}$ where each pair $A_i, B_i \subseteq E(G)$ is an edge bipartition of G and $\beta_i \colon M(A_i,B_i) \to \{1,\dots,|M(A_i,B_i)|\}$ such that $s \geq \log(n)/(r+1)^{32r^2+8r}$, and, for every i, j, $1 \leq i < j \leq s$,*

(i) $|M(A_i,B_i)| = |M(A_j,B_j)| \leq 2r$,
(ii) $A_i \subsetneq A_j$, $B_i \supsetneq B_j$, *and*
(iii) $G\langle A_i\rangle \circ G\langle B_j\rangle$ *is r-outerplanar, where $G\langle A_i\rangle$ is understood to be β_i-boundaried and $G\langle B_j\rangle$ is understood to be β_j-boundaried.*

To gain some intuition for Theorem 2 note that each $M(A_i,B_i)$ is a separator, separating its left side $G\langle A_i\rangle$ from its right side $G\langle B_i\rangle$ in G. Statement (ii) ensures that each left sides contains all previous left sides, that is, the separators are nested. Statement (iii) ensures that for any two separators in the sequence, we can glue their left and right sides and again obtain an r-outerplanar graph. In this new graph, the vertices inbetween the separators are missing—these will be the vertices which are not crucial to obtain an r-outerplanar support.

The reason why we can prove the lower bound on the length of the sequence is basically because r-outerplanar graphs have a tree-like structure, whence large r-outerplanar graphs have a long "path" in this structure, and a long path in such a structure induces many nested separators from which we can glean the separators that are amenable to Statement (iii).

We next formalize the crucial vertices for obtaining an r-outerplanar support. These are the vertices in a smallest representative support, defined as follows.

Definition 2 (Representative support). *Let $\mathcal{H}$ be a hypergraph. A graph G is a* representative support *for $\mathcal{H}$ if $V(G) \subseteq V(\mathcal{H})$, graph G is a support for subhypergraph $\mathcal{H}|_{V(G)}$ shrunken to $V(G)$, and every vertex in $V(\mathcal{H}) \setminus V(G)$ is covered in $\mathcal{H}$ by some vertex in $V(G)$.*

[7] Results labeled by $\star$ are deferred to the Appendix.

Using the sequence of separators from Theorem 2, we show that the size of a smallest representative r-outerplanar support is upper-bounded by a function of m and r. To this end, we take an initial support, find two separators whose vertices in between we can remove and reattach as non-crucial vertices, that is, vertices not in a representative support. Intuitively, the two separators have to have the same "status" with respect to the hyperedges that cross them. We formalize this as follows.

Definition 3 (Edge-bipartition signature). *Let $\mathcal{H} = (V, \mathcal{E})$ be a hypergraph and let G be a representative planar support for $\mathcal{H}$. Let (A, B, β) be a tuple where (A, B) is an edge bipartition of G, and $\beta\colon M(A, B) \to \{1, \ldots, |M(A, B)|\}$. Let $\ell \coloneqq |M(A, B)|$. The* signature *of (A, B, β) is a triple $(\mathcal{T}, \phi, \mathcal{C})$, where*

- *$\mathcal{T} \coloneqq \{[u]_\tau \mid u \in \bigcup A\}$ is the set of twin classes in $\bigcup A$,*
- *$\phi : \{1, \ldots, \ell\} \to \{[u]_\tau \mid u \in V\}\colon j \mapsto [\beta^{-1}(j)]_\tau$ maps each index of a vertex in $M(A, B)$ to the twin class of that vertex, and*
- *$\mathcal{C} \coloneqq \{(F, \gamma_F) \mid F \in \mathcal{E}\}$, where γ_F is the relation on $\{1, \ldots, \ell\}$ defined by $(i, j) \in \gamma_F$ whenever $\beta^{-1}(i), \beta^{-1}(j) \in F$ and $\beta^{-1}(i)$ is connected to $\beta^{-1}(j)$ in $G\langle B\rangle[F \cap \bigcup B]$. Herein, $G\langle B\rangle[F \cap \bigcup B]$ is the subgraph of $G\langle B\rangle$ induced by $F \cap \bigcup B$.*

We have the following upper bound on the number of different separator states.

Lemma 1 ($\star$). *In a sequence $((A_i, B_i, \beta_i))_{i=1}^{s}$ as in Theorem 2 the number of distinct edge-bipartition signatures is upper-bounded by $2^{m\cdot(2r^2+r+1)}$.*

As before, let $\psi(m, r) \coloneqq 2^{6r\cdot 2^{m\cdot(2r^2+r+1)}\cdot(r+1)^{32r^2+8r}}$.

Lemma 2. *If a hypergraph $\mathcal{H} = (V, \mathcal{E})$ has an r-outerplanar support, then it has a representative r-outerplanar support with at most $\psi(m, r)$ vertices.*

Proof. Let $G = (W, E)$ be a representative r-outerplanar support for $\mathcal{H}$ with the minimum number of vertices and fix a corresponding planar embedding. Assume towards a contradiction that $|W| > \psi(m, r)$. We show that there is a representative support for $\mathcal{H}$ with less than $\psi(m, r)$ vertices.

We aim to apply Theorem 2 to G. For this we need that G is connected and does not contain any bridges. Indeed, if G is not connected, then add edges between its connected components in a tree-like fashion. This does not affect the outerplanarity number of G (although it adds bridges). If G has a bridge $\{u, v\}$, then proceed as follows. At least one of the ends of the bridge, say v, has degree at least two because $|W| > \psi(m, r) \geq 2$. One neighbor $w \neq u$ of v is incident with the same face as u, because $\{u, v\}$ is a bridge. Add the edge $\{u, w\}$. Thus, edge $\{u, v\}$ ceases to be a bridge. We can embed $\{u, w\}$ in such a way that the face $\mathfrak{F}$ incident with u, v, and w is split into one face that is incident with only $\{u, v, w\}$ and devoid of any other vertex, and one face $\mathfrak{F}'$ that is incident with all the vertices that are incident with $\mathfrak{F}$ including u, v, and w. This implies that each vertex retains its layer L_i, meaning that G remains r-outerplanar. Thus, we may assume that G is connected, bridgeless, and r-outerplanar.

Since G contains more than $\psi(m,r)$ vertices, there is a sequence $\mathcal{S} = ((A_i, B_i, \beta_i))_{i=1}^{s}$ as in Theorem 2 of length at least

$$s \geq \frac{\log(\psi(m,r))}{(r+1)^{32r^2+8r}} = \frac{6r \cdot 2^{m\cdot(2r^2+r+1)} \cdot (r+1)^{32r^2+8r}}{(r+1)^{32r^2+8r}} = 6r \cdot 2^{m\cdot(2r^2+r+1)}.$$

Since there are less than $2^{m\cdot(2r^2+r+1)}$ different signatures in $\mathcal{S}$ (Lemma 1), there are $6r$ elements of $\mathcal{S}$ with the same signature. Note that each middle set $M(A_i, B_i)$ induces a plane graph in G and, since $|M(A_i, B_i)| \leq 2r$, induces at most $\max\{1, 3|M(A_i, B_i)| - 6\} \leq \max\{1, 6r-6\}$ edges. Thus, there are two edge bipartitions (A_i, B_i, β_i) and (A_j, B_j, β_j), $i < j$, in $\mathcal{S}$ with the same signature such that the middle sets $M(A_i, B_i)$, $M(A_j, B_j)$ differ in at least one vertex.

Let $G_{ij} \coloneqq G\langle A_i\rangle \circ G\langle B_j\rangle$, wherein $G\langle A_i\rangle$ is β_i-boundaried and $G\langle B_j\rangle$ is β_j-boundaried. Let $W' \coloneqq V(G_{ij})$, where we assume that $W' \cap M(A_j, B_j) \subseteq M(A_i, B_i)$ for the sake of a simpler notation. Note that $W \setminus W' \neq \emptyset$ since the middle sets of the two edge bipartitions differ in at least one vertex and since $A_i \subsetneq A_j$.

We prove that G_{ij} is a representative support for $\mathcal{H}$, that is, each vertex $V \setminus W'$ is covered by some vertex in W' in $\mathcal{H}$ and that G_{ij} is a support for $\mathcal{H}|_{W'}$. Since G_{ij} is r-outerplanar by Theorem 2, Statement (iii), this contradicts the choice of G according to the minimum number of vertices, thus proving the lemma.

To prove that each vertex $V \setminus W'$ is covered by some vertex in W', we show that $\{[u]_\tau \mid u \in V\} = \{[u]_\tau \mid u \in W'\}$. Since $G = (W, E)$ is a representative support, $\{[u]_\tau \mid u \in V\} = \{[u]_\tau \mid u \in W\}$. Furthermore, by the definition of signature, we have $\{[u]_\tau \mid u \in \bigcup A_i\} = \{[u]_\tau \mid u \in \bigcup A_j\}$. Thus, for each vertex $u \in W \setminus W'$, there is a vertex $v \in W'$ with $[u]_\tau = [v]_\tau$, meaning that, indeed, $\{[u]_\tau \mid u \in V\} = \{[u]_\tau \mid u \in W'\}$.

To show that G_{ij} is a representative support it remains to show that it is a support for $\mathcal{H}|_{W'}$, that is, each hyperedge F' of $\mathcal{H}|_{W'}$ induces a connected graph $G_{ij}[F']$. Let F be a hyperedge of $\mathcal{H}$ such that $F \cap W' = F'$. Observe that such a hyperedge F exists and that $G[F \cap W]$ is connected since G is a representative support of $\mathcal{H}$. Denote by S_k the middle set $M(A_k, B_k)$ of (A_k, B_k) in G for $k \in \{i, j\}$ and by S the middle set $M(A_i, B_j) = S_i = S_j$ of (A_i, B_j) in G_{ij}.

To show that $G_{ij}[F']$ is connected, consider first the case that $F \cap (S_i \cup S_j) = \emptyset$. Since each vertex in $V \setminus W'$ is covered by a vertex in W' we have that each vertex in F is contained in either $G\langle A_i\rangle$ or $G\langle B_j\rangle$ along with all edges of $G[F]$. All these edges are also present in G_{ij} whence $G_{ij}[F']$ is connected.

Now consider the case that $F \cap (S_i \cup S_j) \neq \emptyset$. Since S_i and S_j are separators in G, each vertex in $F \setminus (S_i \cup S_j)$ is connected in $G[F]$ to some vertex in S_i or S_j via a path with internal vertices in $F \setminus (S_i \cup S_j)$. We consider the connectivity relation of their corresponding vertices in S. To this end, for a graph H and $T \subseteq V(H)$ use $\gamma(T, H)$ for the equivalence relation on T of connectivity in H. That is, for $u, v \in T$ we have $(u, v) \in \gamma(T, H)$ if u and v are connected in H. Using this terminology, since both S_i and S_j equal S in G_{ij}, to show that $G_{ij}[F']$ is connected, it is enough to prove that the transitive closure δ of $\gamma(F' \cap S, G_{ij}\langle A_i\rangle) \cup \gamma(F' \cap S, G_{ij}\langle B_j\rangle)$ contains only one equivalence class.

Denote by $\hat{G}$ the graph obtained from G by identifying each $v \in S_i$ with $\beta_j^{-1}(\beta_i(v)) \in S_j$, hence, identifying S_i and S_j, resulting in the set S. Relation $\alpha := \gamma(F \cap S, \hat{G})$ has only one equivalence class and, moreover, it is the transitive closure of $\gamma(F \cap S_i, G\langle A_i \rangle) \cup \gamma(F \cap S, \hat{G}\langle B_i \setminus B_j \rangle) \cup \gamma(F \cap S_j, G\langle B_j \rangle)$, wherein we identify each $v \in S_i$ with $\beta_j^{-1}(\beta_i(v)) \in S_j$ as above and, thus, $S_i = S_j = S$. We have $\gamma(F' \cap S, G_{ij}\langle A_i \rangle) = \gamma(F \cap S_i, G\langle A_i \rangle)$ and $\gamma(F' \cap S, G_{ij}\langle B_j \rangle) = \gamma(F \cap S_j, G\langle B_j \rangle)$. Thus for $\alpha = \delta$ it suffices to prove that $\gamma(F \cap S, \hat{G}\langle B_i \setminus B_j \rangle) \subseteq \gamma(F' \cap S_j, G_{ij}\langle B_j \rangle)$. Indeed, the left-hand side $\gamma(F \cap S, \hat{G}\langle B_i \setminus B_j \rangle)$ is contained in $\gamma(F \cap S_i, G\langle B_i \rangle)$. Let $(\mathcal{T}, \phi, \mathcal{C})$ be the signature of (A_i, B_i, β_i) and (A_j, B_j, β_j) and $(F, \gamma_F) \in \mathcal{C}$. Note that $\gamma(F \cap S_i, G\langle B_i \rangle) = \gamma_F = \gamma(F \cap S_j, G\langle B_j \rangle)$ where we abuse notation and set $u = \beta_i(u)$ for $u \in S_i$ and $v = \beta_j(v)$ for $v \in S_j$. Hence, $\gamma(F \cap S, \hat{G}\langle B_i \setminus B_j \rangle) \subseteq \gamma(F \cap S_j, G\langle B_j \rangle) = \gamma(F' \cap S_j, G\langle B_j \rangle) = \gamma(F' \cap S_j, G_{ij}\langle B_i \rangle)$. Thus, indeed, $\delta = \alpha$, that is, F' is connected. □

We now use the upper bound on the number of vertices in representative supports to get rid of superfluous twins. First, we show that representative supports can be extended to obtain a support.

Lemma 3. *Let $G = (W, E)$ be a representative r-outerplanar support for a hypergraph $\mathcal{H} = (V, \mathcal{E})$. Then $\mathcal{H}$ has an r-outerplanar support in which all vertices of $V \setminus W$ have degree one.*

Proof. Let G' be the graph obtained from G by making each vertex v of $V \setminus W$ a degree-one neighbor of a vertex in W that covers v (such a vertex exists by the definition of representative support). Clearly, the resulting graph is planar. It is also r-outerplanar, which can be seen by adapting an r-outerplanar embedding of G for G': If the neighbor v of a new degree-one vertex u is in L_1, then place u in the outer face. If $v \in L_i$, $i > 1$, then place u in a face which is incident with v and a vertex in L_{i-1} (such a face exists by the definition of L_i).

It remains to show that G' is a support for $\mathcal{H}$. Consider a hyperedge $F \in \mathcal{E}$. Since G is a representative support for $\mathcal{H}$, we have that $F \cap W$ is nonempty and that $G[F \cap W]$ is connected. In G', each vertex $u \in F \setminus W$ is adjacent to some vertex $v \in W$ that covers u. Hence $v \in F$. Thus, $G'[F]$ is connected as $G'[F \cap W]$ is connected and all vertices in $F \setminus W$ are neighbors of a vertex in $F \cap W$. □

We now use Lemma 3 to show that, if there is a twin class that contains more vertices than a small representative support, then we can safely remove one vertex from this twin class.

Lemma 4. *Let $\ell \in \mathbb{N}$, let $\mathcal{H}$ be a hypergraph, and let $v \in V(\mathcal{H})$ be a vertex such that $|[v]_\tau| \geq \ell$. If $\mathcal{H}$ has a representative r-outerplanar support with less than ℓ vertices, then $\mathcal{H} - v$ has an r-outerplanar support.*

Proof. Let $G = (W, E)$ be a representative r-outerplanar support for $\mathcal{H}$ such that $|W| < \ell$. Then at least one vertex of $[v]_\tau$ is not in W and we can assume that this vertex is v without loss of generality. Thus, $\mathcal{H}$ has an r-outerplanar support G' in which v has degree one by Lemma 3. The graph $G' - v$ is an r-outerplanar support for $\mathcal{H} - v$: For each hyperedge F in $\mathcal{H} - v$, we have that $G'[F \setminus \{v\}]$ is connected because v is not a cut-vertex in $G'[F]$ (since it has degree one). □

Now we combine the observations above with the fact that there are small r-outerplanar supports to prove that Rule 1 is correct.

Proof (Correctness of Rule 1). Consider an instance $\mathcal{H} = (V, \mathcal{E})$ of r-Outerplanar Support to which Rule 1 is applicable and let $v \in V$ be a vertex to be removed, that is, v is contained in a twin class of size more than $\psi(m, r)$. By Lemma 2, if $\mathcal{H}$ has an r-outerplanar support, then it has a representative r-outerplanar support with at most $\psi(m, r)$ vertices. By Lemma 4, this implies that $\mathcal{H} - v$ has an r-outerplanar support. Moreover, if $\mathcal{H} - v$ has an r-outerplanar support, then this r-outerplanar support is a representative r-outerplanar support for $\mathcal{H}$. By Lemma 3, this implies that $\mathcal{H}$ has an r-outerplanar support. Therefore, $\mathcal{H}$ and $\mathcal{H} - v$ are equivalent instances, and v can be safely removed from $\mathcal{H}$. □

5 Concluding remarks

The main contribution of this work is to show that twins may be crucial for instances of r-Outerplanar Support but the number of crucial twins is upper-bounded in terms of the number m of hyperedges and the outerplanarity number r of a support. As a result, we can safely remove non-crucial twins. More specifically, in linear time we can transform any instance of r-Outerplanar Support into an equivalent one whose size is upper-bounded by a function of m and r only. In turn, this implies fixed-parameter tractability with respect to $m+r$. It is fair to say, however, that due to the strong exponential growth in m and r this result is mainly of classification nature. Improved bounds (perhaps based on further data reduction rules) are highly desirable for practical applications.

Two further directions for future research are as follows. First, above we only showed how to reduce the size of the input instance. We also need an efficient algorithm to construct an r-outerplanar support for such an instance. As a first step, it would be interesting to improve on the $n^{\mathrm{O}(n)}$-time brute-force algorithm that simply enumerates all n-vertex planar graphs and tests whether one of them is an r-outerplanar support.[8]

Second, it is interesting to gear the parameters under consideration more towards practice. In Section 4 above we attached signatures to each edge bipartition in a sequence of edge bipartitions of a support and we could reduce our input only if there were sufficiently many edge bipartitions with the same signature. This signature contained, among other information, the twin class of each vertex of the separator induced by the edge bipartition. Clearly, if all of these at least 2^{mr} different types of signatures are present, this will lead to an illegible drawing of the hypergraph (and still, in absence of better upper bounds, we cannot reduce our input). It seems thus worthwhile to contemplate parameters that capture legibility of the hypergraph drawing by restricting further the number of possible signatures.

[8] Each planar graph has an ordering of the vertices such that each vertex has at most five neighbors later in the ordering. To achieve $n^{\mathrm{O}(n)}$ enumeration time we simply guess such an ordering and then for each vertex its at most five later neighbors.

Finally, an obvious open question is whether finding a *planar* support is (linear-time) fixed-parameter tractable with respect to the number m of hyperedges only. A promising direction might be to show that there is a planar representative support (as in Definition 2) which has treewidth upper-bounded by a function of m. From this, we would get a sequence of gluable subgraphs similarly to the one we have used here, amenable to the same approach as in Section 4.

Bibliography

[1] B. Alsallakh, L. Micallef, W. Aigner, H. Hauser, S. Miksch, and P. J. Rodgers. The state-of-the-art of set visualization. *Computer Graphics Forum*, 35(1): 234–260, 2016.

[2] C. Beeri, R. Fagin, D. Maier, and M. Yannakakis. On the desirability of acyclic database schemes. *Journal of the ACM*, 30(3):479–513, 1983.

[3] T. Biedl. On triangulating k-outerplanar graphs. *Discrete Applied Mathematics*, 181:275–279, 2015.

[4] R. E. Bixby and D. K. Wagner. An almost linear-time algorithm for graph realization. *Mathematics of Operations Research*, 13(1):99–123, 1988.

[5] U. Brandes, S. Cornelsen, B. Pampel, and A. Sallaberry. Blocks of hypergraphs—Applied to hypergraphs and outerplanarity. In C. S. Iliopoulos and W. F. Smyth, editors, *Proceedings of the 21st International Workshop on Combinatorial Algorithms (IWOCA '10)*, volume 6460 of *LNCS*, pages 201–211. Springer Berlin Heidelberg, 2011.

[6] U. Brandes, S. Cornelsen, B. Pampel, and A. Sallaberry. Path-based supports for hypergraphs. *Journal of Discrete Algorithms*, 14:248–261, 2012.

[7] K. Buchin, M. J. van Kreveld, H. Meijer, B. Speckmann, and K. Verbeek. On planar supports for hypergraphs. *Journal of Graph Algorithms and Applications*, 15(4):533–549, 2011.

[8] J. Chen, C. Komusiewicz, R. Niedermeier, M. Sorge, O. Suchý, and M. Weller. Polynomial-time data reduction for the subset interconnection design problem. *SIAM Journal on Discrete Mathematics*, 29(1):1–25, 2015.

[9] M. Cygan, F. V. Fomin, L. Kowalik, D. Lokshtanov, D. Marx, M. Pilipczuk, M. Pilipczuk, and S. Saurabh. *Parameterized Algorithms*. Springer, 2015.

[10] L. E. Dickson. Finiteness of the odd perfect and primitive abundant numbers with n distinct prime factors. *American Journal of Mathematics*, 35(4): 413–422, 1913.

[11] R. Diestel. *Graph Theory*, volume 173 of *Graduate Texts in Mathematics*. Springer, 4th edition edition, 2010.

[12] K. Dinkla, M. J. van Kreveld, B. Speckmann, and M. A. Westenberg. Kelp diagrams: Point set membership visualization. *Computer Graphics Forum*, 31(3):875–884, 2012.

[13] F. Dorn, E. Penninkx, H. L. Bodlaender, and F. V. Fomin. Efficient exact algorithms on planar graphs: Exploiting sphere cut decompositions. *Algorithmica*, 58(3):790–810, 2010.

[14] T. Eschbach, W. Günther, and B. Becker. Orthogonal hypergraph drawing for improved visibility. *Journal of Graph Algorithms and Applications*, 10 (2):141–157, 2006.

[15] J. Flower, A. Fish, and J. Howse. Euler diagram generation. *Journal of Visual Languages & Computing*, 19(6):675–694, 2008.

[16] Q. Gu and H. Tamaki. Optimal branch-decomposition of planar graphs in $o(n^3)$ time. *ACM Transactions on Algorithms*, 4(3), 2008.

[17] M. Habib, C. Paul, and L. Viennot. Partition refinement techniques: An interesting algorithmic tool kit. *International Journal of Foundations of Computer Science*, 10(2):147–170, 1999.

[18] D. S. Johnson and H. O. Pollak. Hypergraph planarity and the complexity of drawing Venn diagrams. *Journal of Graph Theory*, 11(3):309–325, 1987.

[19] M. Kaufmann, M. J. van Kreveld, and B. Speckmann. Subdivision drawings of hypergraphs. In I. G. Tollis and M. Patrignani, editors, *Proceedings of the 16th International Symposium on Graph Drawing (GD '08)*, volume 5417 of *LNCS*, pages 396–407. Springer Berlin Heidelberg, 2008.

[20] B. Klemz, T. Mchedlidze, and M. Nöllenburg. Minimum tree supports for hypergraphs and low-concurrency euler diagrams. In R. Ravi and I. L. Gørtz, editors, *Proceedings of the 14th Scandinavian Symposium on Algorithm Theory (SWAT '14)*, volume 8503 of *LNCS*, pages 265–276. Springer International Publishing, 2014.

[21] E. Korach and M. Stern. The clustering matroid and the optimal clustering tree. *Mathematical Programming*, 98(1-3):385–414, 2003.

[22] D. Marx and M. Pilipczuk. Optimal parameterized algorithms for planar facility location problems using voronoi diagrams. *CoRR*, abs/1504.05476, 2015. URL http://arxiv.org/abs/1504.05476.

[23] J. Munkres. *Topology*. Prentice Hall, 2nd edition edition, 2000.

[24] E. Mäkinen. How to draw a hypergraph. *International Journal of Computer Mathematics*, 34:178–185, 1990.

[25] P. D. Seymour and R. Thomas. Call routing and the ratcatcher. *Combinatorica*, 14(2):217–241, 1994.

[26] R. Tarjan and M. Yannakakis. Simple linear-time algorithms to test chordality of graphs, test acyclicity of hypergraphs, and selectively reduce acyclic hypergraphs. *SIAM Journal on Computing*, 13(3):566–579, 1984.

[27] W. T. Tutte. How to draw a graph. *Proceedings of the London Mathematical Society*, 13(3):743–768, 1963.

[28] A. Verroust and M. Viaud. Ensuring the drawability of extended euler diagrams for up to 8 sets. In A. F. Blackwell, K. Marriott, and A. Shimojima, editors, *Proceedings of the Third International Conference on Diagrammatic Representation and Inference (Diagrams '04)*, volume 2980 of *LNCS*, pages 128–141. Springer Berlin Heidelberg, 2004.

Appendix

A Omitted content from Section 3

By introducing several copies of the hypergraph constructed in Section 3, we enforce that each vertex of a twin class is contained in a distinct face of any planar embedding of the support.

Fix an integer $\ell \in \mathbb{N}$. To construct the hypergraph $\mathcal{H}$, copy the vertex set V from Section 3 ℓ times, and let $V_i := \{a_i, b_i, c_i, d_i, v_{i,a}, v_{i,b}, v_{i,d}, u_{i,b}, u_{i,c}, u_{i,d}, t_i, t_i'\}$ denote the vertex set of the ith copy. Within each copy, add the size-two solid hyperedges as in Section 3. Then, add a distinguished vertex v^*, and add the size-two hyperedges $\{a_i, v^*\}$, $\{b_i, v^*\}$, and $\{c_i, v^*\}$ to $\mathcal{H}$ for each $i \in \{1, \ldots, \ell\}$. Vertex v^* serves as a conduit to connect subgraphs induced by hyperedges that contain vertices from all the copies, which we are about to introduce.

Let $X = \{x \mid i \in \{1, \ldots, \ell\}\}$ for each $(X, x) \in \{(A, a_i), (B, b_i), (C, c_i), (V_a, v_{i,a}), (V_b, v_{i,b}), (U_b, u_{i,b}), (U_c, u_{i,c})\}$, and let $T := \{t_i, t_i' \mid i \in \{1, \ldots, \ell\}\}$. The final hyperedges in $\mathcal{H}$ are

$$A \cup C \cup V_a \cup T \cup \{v^*\}, \qquad A \cup C \cup V_b \cup T \cup \{v^*\},$$
$$B \cup C \cup V_a \cup T \cup \{v^*\}, \qquad B \cup C \cup V_b \cup T \cup \{v^*\},$$
$$B \cup A \cup U_b \cup T \cup \{v^*\}, \qquad B \cup A \cup U_c \cup T \cup \{v^*\},$$
$$C \cup A \cup U_b \cup T \cup \{v^*\}, \quad \text{and } C \cup A \cup U_c \cup T \cup \{v^*\}.$$

Note that T forms a twin class in the resulting hypergraph $\mathcal{H}$. Hypergraph $\mathcal{H}$ has a planar support because v^* can be used to connect for each $\mathcal{H}[V_i]$ the partial supports that are obtained by copying the support for the simple example from Section 3.

We claim that, for each $t \in T$, hypergraph $\mathcal{H} - t$ does not have a planar support. To see this, assume that there is a support G and consider the planar embedding of $G - T$. Note that, in this embedding, no vertex $v_{a,i}, v_{b,i} \in V_a \cup V_b$ is in the same face as any vertex $u_{b,i}, u_{c,i} \in U_b \cup U_c$, because these vertices are in two connected components incident with d_i, a_i, and b_i, or d_i, b_i, and c_i, respectively, as in the simple example before. In addition, no vertex in $V_a \cup V_b \cup U_b \cup U_c$ from the ith copy is incident with the same face as a vertex in $V_a \cup V_b \cup U_b \cup U_c$ from the jth copy, where $i \neq j$. This is because all the vertices in $A \cup B \cup C$ have to be incident with the same face, due to the connections of v^* to each copy of a, b, and c. Thus, each pair of copies of $\{a, b, v_a, v_b, v_d\}$ or $\{a, b, u_b, u_c, u_d\}$ defines one of 2ℓ distinct faces. As before, each of these faces has to contain a vertex from T, thus, if $|T| < 2\ell$, then there is no planar support.

B Existence of a number of important twins

In this section we show the following.

Theorem 3. *There exists a function $\psi\colon \mathbb{N} \to \mathbb{N}$ with the following property. For each $m \in \mathbb{N}$ and every hypergraph $\mathcal{H}$ that has m hyperedges, out of each twin class of $\mathcal{H}$, we can remove all but $\psi(m)$ arbitrary twins such that the resulting hypergraph has a planar support if and only if $\mathcal{H}$ has one.*

The basic idea is that adding twins is not detrimental. If we have a planar support for $\mathcal{H}$, then we can make a new twin adjacent to one of its already present twins, so that the resulting graph remains planar. Reversing this idea, from each hypergraph with a planar support, by deleting twins we can obtain a *minimal* hypergraph $\mathcal{H}'$ which also has a planar support but from which no further twins can be deleted while maintaining the property of having a planar support. Using Dickson's lemma (see below for details) it is not hard to show that there is a function ϕ such that, for each fixed number m of hyperedges, there are only $\phi(m)$ such minimal hypergraphs. (A priori, we do not know of a way to compute ϕ, however. A way to compute ϕ in the case of r-outerplanar supports arises from our results in Section 4.) Clearly, among these minimal hypergraphs, one has a largest twin class, whose size we can set to $\psi(|\mathcal{E}(\mathcal{H})|)$.

We now formalize the above approach. Denote by $\mathbb{S}$ the set of hypergraphs which have a planar support. As mentioned, $\mathbb{S}$ is *closed under adding twins*, that is taking an arbitrary hypergraph $\mathbb{S}$ and adding a twin to it yields another hypergraph in $\mathbb{S}$.

Proof (Theorem 3). We first define a quasi-order $\preceq$ on the family of hypergraphs with m hyperedges. (A *quasi-order* is reflexive and transitive.) To define $\preceq$, we say that $\mathcal{H} \preceq \mathcal{G}$ if $\mathcal{H}$ can be obtained from $\mathcal{G}$ by iteratively removing a vertex that has a twin. If we allow zero removals so that $\preceq$ is reflexive, it is clear that $\preceq$ is a quasi-order. Moreover, if $\mathcal{H} \in \mathbb{S}$ and $\mathcal{H} \preceq \mathcal{G}$, then $\mathcal{G} \in \mathbb{S}$ since $\mathbb{S}$ is closed under adding twins.

Next we show that, for every $m \in \mathbb{N}$, the family $\mathbb{F}_m$ of hypergraphs in $\mathbb{S}$ that are minimal under $\preceq$ is finite. Consider the representation of an m-hyperedge hypergraph $\mathcal{H}$ as a 2^m-tuple $t_{\mathcal{H}} \in \mathbb{N}^{2^m}$, each entry of which represents the size of a distinct twin class. The set of such tuples is quasi-ordered by a natural extension of $\leq$, namely $(a_1, \ldots, a_\ell) \leq (b_1, \ldots, b_\ell)$ if $a_i \leq b_i$ for each $i \in \{1, \ldots, \ell\}$. We now lead the assumption that $\mathbb{F}_m$ is infinite to a contradiction by using Dickson's Lemma [10, Lemma A].

If $\mathbb{F}_m$ is infinite, then there is an infinite subset $\mathbb{F}'_m$ of hypergraphs which have the same (nonempty) twin classes. For hypergraphs $\mathcal{H}, \mathcal{G}$ with the same twin classes, $t_{\mathcal{H}} \leq t_{\mathcal{G}}$ implies $\mathcal{H} \preceq \mathcal{G}$. Thus, $\mathbb{F}'_m$ gives an infinite set T of tuples that are pairwise incomparable under $\leq$. Dickson's Lemma states that for every set $S \subseteq \mathbb{N}^\ell$ there exists a finite subset $S' \subseteq S$ such that for each $s \in S$ there is an $s' \in S'$ with $s' \leq s$. This is a contradiction to T containing infinitely many incomparable tuples. Hence, $\mathbb{F}_m$ is finite.

Finally, to obtain the function ψ in the theorem, we set $\psi(m)$ as the largest size of a twin class of a hypergraph in $\mathbb{F}_m$. □

We emphasize that the proof of Theorem 3 does not provide insight into why it is that a hypergraph without a planar support suddenly gets one when

adding a twin. That is, we currently do not know how to compute the set $\mathbb{F}_m$ of minimal yes-instances, making this result inapplicable at the moment. Hence, to eventually obtain implementable algorithms that are able to deal with any input, it is important to constructivize Theorem 3. This we do for r-outerplanar supports in Section 4.

Acknowledgments. We thank an anonymous referee who pointed out the ideas for the proof in Appendix B.

C Additional Preliminaries

For the proof of Theorem 2 in Appendix D we need some additional preliminaries. In particular, we need some notions from topology to more precisely argue about embeddings, and the notion of a face path in a plane graph and (sphere-cut) branch decompositions as central tools in the proof.

Topology. A *topological space* is a tuple $\mathfrak{X} = (X, \mathcal{F})$ of a set X, called *universe*, and a collection $\mathcal{F}$ of subsets of X, called *topology*, that satisfy the following properties:

- The empty set $\emptyset$ and X are in $\mathcal{F}$.
- The union of the elements of any subcollection of $\mathcal{F}$ is in $\mathcal{F}$.
- The intersection of the elements of any finite subcollection of $\mathcal{F}$ is in $\mathcal{F}$.

Each set in $\mathcal{F}$ is called *open*. A *closed set* is the complement of an open set. (The empty set and X are both open and closed.)

We consider here the topological space $\mathfrak{R}^\ell = (\mathbb{R}^\ell, \mathcal{F})$ where $\mathcal{F}$ is the standard topology of $\mathbb{R}^\ell$, that is, $\mathcal{F}$ is the closure under union and finite intersection of the open balls $\{\boldsymbol{x} \in \mathbb{R}^\ell \mid \|\boldsymbol{x} - \boldsymbol{y}\| < d\}$ for $d \in \mathbb{R}$, $\boldsymbol{y} \in \mathbb{R}^\ell$, where $\|\cdot\|$ is the Euclidean norm.

A *topological subspace* $\mathfrak{T} \subseteq \mathfrak{S}$ of a topological space $\mathfrak{S}$ is a topological space whose universe is a subset of the universe of $\mathfrak{S}$. We always assume topological subspaces to carry the *subspace topology*, that is, the open sets of $\mathfrak{T}$ are the intersections of the open sets of $\mathfrak{S}$ with the universe of $\mathfrak{T}$. We also say that $\mathfrak{T}$ is the topological subspace *induced* by the universe of $\mathfrak{T}$.

Important topological subspaces of $\mathfrak{R}^\ell$ are, with a slight abuse of notation,

- the *plane* $\mathfrak{R}^2$,
- the *sphere*, whose universe is $\{(x, y, z) \in \mathbb{R}^3 \mid x^2 + y^2 + z^2 = 1\}$,
- the *closed disk*, whose universe is $\{(x, y) \in \mathbb{R}^2 \mid x^2 + y^2 \leq 1\}$,
- the *open disk*, whose universe is $\{(x, y) \in \mathbb{R}^2 \mid x^2 + y^2 < 1\}$, and
- the *circle*, whose universe is $\{(x, y) \in \mathbb{R}^2 \mid x^2 + y^2 = 1\}$.

A *homeomorphism* ϕ between two topological spaces is a bijection ϕ between the two corresponding universes such that both ϕ and ϕ^{-1} are continuous. We often refer to a *subspace* $\mathfrak{X}$ *in a topological space* $\mathfrak{Y}$ (for example, a circle in a plane), by which we mean a topological subspace of $\mathfrak{Y}$ which is homeomorphic to $\mathfrak{X}$.

An *arc* is a topological space that is homeomorphic to the closed interval $[0, 1] \subseteq \mathfrak{R}^1$. The images of 0 and 1 under a corresponding homeomorphism are the

endpoints of the arc, which *links* them and runs *between* them. Let $\mathfrak{X} = (X, \mathcal{X})$ be a topological space. Being linked by an arc in $\mathfrak{X}$ defines an equivalence relation on X. The topological subspaces induced by the equivalence classes of this relation are called *regions*. We say that a closed set C in a topological space $\mathfrak{S}$ *separates* $\mathfrak{S}$ into the regions of the subspace of $\mathfrak{S}$ induced by $S \setminus C$ where S is the universe of $\mathfrak{S}$.

For more on topology, see Munkres [23], for example.

Embeddings of graphs into the plane and sphere. An *embedding* of a graph $G = (V, E)$ into the plane $\mathfrak{R}^2$ (into the sphere $\mathfrak{S}$) is a tuple $(\mathfrak{V}, \mathcal{E})$ and a bijection $\phi\colon V \to \mathfrak{V}$ such that

- $\mathfrak{V} \subseteq \mathfrak{R}^2$ ($\mathfrak{V} \subseteq \mathfrak{S}$),
- $\mathcal{E}$ is a set of arcs in $\mathfrak{R}^2$ (in $\mathfrak{S}$) with endpoints in $\mathfrak{V}$,
- the interior of any arc in $\mathcal{E}$ (that is, the arc without its endpoints) contains no point in $\mathfrak{V}$ and no point of any other arc in $\mathcal{E}$, and
- $u, v \in V$ are adjacent in G if and only if $\phi(u)$ is linked to $\phi(v)$ by an arc in $\mathcal{E}$.

The regions in $\mathfrak{R}^2 \setminus (\bigcup \mathcal{E})$ (in $\mathfrak{S} \setminus (\bigcup \mathcal{E})$) are called *faces*.

A *planar graph* is a graph which has an embedding in the plane or, equivalently, in the sphere. A *plane graph* $G = (V, E)$ is a planar graph given with a fixed embedding in the plane. An $\mathfrak{S}$*-plane graph* G is a planar graph given with a fixed embedding in the sphere. For notational convenience, we refer to the sets V and $\mathfrak{V}$ as well as E and $\mathcal{E}$ interchangeably. Moreover, we sometimes identify G with the set of points $\mathfrak{V} \cup \bigcup \mathcal{E}$.

A *noose in an* $\mathfrak{S}$*-plane graph* G is a circle in $\mathfrak{S}$ whose intersection with G is contained in $V(G)$. Note that every noose separates $\mathfrak{S}$ into two open disks.

Layer decompositions, outerplanar graphs, face paths. The face of unbounded size in the embedding of a plane graph G is called *outer face*. The *layer decomposition* of G with respect to the embedding is a partition of V into layers $L_1 \uplus \cdots \uplus L_r$ and is defined inductively as follows. Layer L_1 is the set of vertices that lie on the outer face of G, and, for each $i \in \{2, \ldots, r\}$, layer L_i is the set of vertices that lie on the outer face of $G - (\bigcup_{j=1}^{i-1} L_j)$. The graph G is called r*-outerplanar* if it has an embedding with a layer decomposition consisting of at most r layers. We denote by the *outerplanarity number* of G the minimum r such that G is r-outerplanar. If $r = 1$, then G is simply said to be *outerplanar*. A *face path* is an alternating sequence of faces and vertices such that two consecutive elements are incident with one another. The first and last element of a face path are called its *ends*. Note that the ends of a face path may be two vertices, two faces, or a face and a vertex. The *length* of a face path is the number of faces in the sequence. Note that a vertex v in layer L_i has a face path of length i from v to the outer face. Moreover, a graph is r-outerplanar if and only if each vertex has a face path of length at most r to the outer face.

Branch decompositions. A *branch decomposition of a graph* G is a tuple (T, λ) where T is a ternary tree, that is, each internal vertex has degree three, and λ is a bijection between the leaves of T and $E(G)$. Every edge $e \in E(T)$ defines

a bipartition of $E(G)$ into A_e, B_e corresponding to the leaves in the connected components of $T - e$. Define the *middle set* $M(e)$ of an edge $e \in E(T)$ to be the set of vertices in G which are incident with both an edge in A_e and B_e. That is,

$$M(e) \coloneqq \{v \in V(G) \mid \exists a \in A_e \exists b \in B_e \colon v \in a \cap b\}.$$

The *width of an edge* $e \in E(T)$ is $|M(e)|$ and the *width of a branch decomposition* (T, λ) is the largest width of an edge in T. The *branchwidth of a graph* G is the smallest width of a branch decomposition of G.

A *sphere-cut branch decomposition of an* $\mathfrak{S}$*-plane graph* G is a branch decomposition (T, λ) of G fulfilling the following additional condition. For every edge $e \in E(T)$, there is a noose $\mathfrak{N}_e$ whose intersection with G is precisely $M(e)$ and, furthermore, the open disks $\mathfrak{D}_1, \mathfrak{D}_2$ into which the noose $\mathfrak{N}_e$ separates $\mathfrak{S}$, can be indexed in such a way that $\mathfrak{D}_1 \cap G = A_e \setminus M(e)$ and $\mathfrak{D}_2 \cap G = B_e \setminus M(e)$. We use the following theorem.

Theorem 4 ([25, 13, 22]). *Let G be a connected, bridgeless, $\mathfrak{S}$-plane graph of branchwidth at most b. There exists a sphere-cut branch decomposition for G of width at most b.*

Dorn et al. [13] first noted that Seymour and Thomas [25] implicitly proved a variant of Theorem 4 in which G is required to have no degree-one vertices rather than no bridges. Marx and Pilipczuk [22] observed a flaw in Dorn et al.'s derivation, showing that bridgelessness is required (and sufficient). The sphere-cut branch decomposition in Theorem 4 can be computed in $O(|V(G)|^3)$ time (see Gu and Tamaki [16]), but we do not need to explicitly construct it below.

D A sequence of gluable edge bipartitions

In this section, given an r-outerplanar graph, we provide a sequence of separators, each of size at most $2r$, that has the following properties.

- Each of the separators separates the graph into a well-defined *left* and *right* side (we say below that this is the left and right side of the separator).
- The separators are *nested*, meaning that each left side of a separator contains all left sides of separators with smaller index in the sequence.
- For every two separators S_i, S_j with $j > i$, gluing the left side of S_i with the right side of S_j yields an r-outerplanar graph. Gluing means to pairwise identify the vertices of S_i and S_j, in particular, $|S_i| = |S_j|$.

We lower-bound the length of the sequence in terms of the number n of vertices in G and the outerplanarity number r.

In Section 4 we then use such a sequence to show that, if a hypergraph is large, then it contains vertices which are not crucial for having an r-outerplanar support. Essentially, we take an r-outerplanar support for the hypergraph, construct the sequence of separators, and use two separators S_1, S_2 in the sequence to show that we can glue the left side of S_1 and the right side of S_2, that is, remove all vertices not contained in the left and right sides and reattach them as (non-crucial)

degree-one vertices. Furthermore, we can do this in such a way that we again get an r-outerplanar support.

To formally define the sequence of separators, we use the following notation. Although the intuition about separators is instructive, it is more convenient to define our sequence of separators in terms of edge bipartitions.

Definition 4. *For an edge bipartition $A, B \subseteq E(G)$ of a graph G, let $M(A, B)$ be the set of vertices in G which are adjacent with both an edge in A and in B, that is,*

$$M(A, B) \coloneqq \{v \in V(G) \mid \exists a \in A \exists b \in B \colon v \in a \cap b\}.$$

We call $M(A, B)$ the middle set *of A, B, similarly to the middle sets in branch decompositions. For an edge set $A \subseteq E(G)$, denote by $G\langle A\rangle \coloneqq (\bigcup_{e \in A} e, A)$ the subgraph induced by A.*

Recall also from Section 2 the definitions of graph gluing, boundary, and boundary labeling.

We prove the following theorem.

Theorem 2. *For every connected, bridgeless, r-outerplanar graph G with n vertices there is a sequence $((A_i, B_i, \beta_i))_{i=1}^{s}$ where $A_i, B_i \subseteq E(G)$ and $\beta_i \colon M(A_i, B_i) \to \{1, \ldots, |M(A_i, B_i)|\}$ such that $s \geq \log(n)/(r+1)^{32r^2+8r}$, and, for every i, j, $1 \leq i < j \leq t$,*

(i) $|M(A_i, B_i)| = |M(A_j, B_j)| \leq 2r$,
(ii) $A_i \subsetneq A_j$, $B_i \supsetneq B_j$, *and*
(iii) $G\langle A_i\rangle \circ G\langle B_j\rangle$ *is r-outerplanar, where $G\langle A_i\rangle$ is understood to be β_i-boundaried and $G\langle B_j\rangle$ is understood to be β_j-boundaried.*

The proof relies crucially on sphere-cut branch decompositions [13, 22]. A *sphere-cut branch decomposition* is a ternary tree T whose leaves one-to-one correspond to the edges of the graph G embedded in the sphere (without edge crossings) that fulfills the following property. For each edge e in T, there is a circle in the sphere that meets G in precisely the middle set of the edge bipartition (A, B) of G induced by the connected components of $T - e$, and moreover, that circle cuts the sphere into two disks such that one of the disks contains only edges from A and the other only from B. Such a circle is also called *noose*. For the precise definitions, see Section 2.

Outline of the proof of Theorem 2. We first transform the planar embedding of G into an embedding in the sphere and apply Theorem 4 from which we obtain a sphere-cut branch decomposition for G of width at most $2r$. The edge bipartitions in Theorem 2 are defined based on the edges in a longest path in the decomposition tree corresponding to the sphere-cut branch decomposition. The longest path in the decomposition tree has length at least $2\log(n)$, and the edges on this path will define a sequence of edge bipartitions, a supersequence of the one in Theorem 2. We define a signature for each bipartition, a string containing $(32r^2+8r)\cdot\log(r+1)+1$ bits, which determines the pairs of edge bipartitions that can be glued so that we

obtain an r-outerplanar graph. The sequence in Theorem 2 is then obtained from those bipartitions which have the same signature. The sphere-cut property of the branch decomposition gives one noose in the sphere for each edge bipartition in the sequence, such that it separates the parts in the edge bipartition from one another. The nooses of the sphere-cut branch decomposition will be crucial in the proof of Statement (iii) in Theorem 2, that is, the r-outerplanarity of the glued graphs.

Let us give some more details concerning the r-outerplanarity of the glued graphs. After sanitizing the nooses, we can assume that they separate the sphere into *left* disks and *right* disks in such a way that each left disk contains all left disks with smaller index. Hence, for each pair of nooses, we can cut out a left disk and a right disk, and glue them along their corresponding nooses such that we again get a sphere. Alongside the sphere, we get a graph embedded in it that corresponds to the left and right sides of the separators induced by the nooses. It then remains to make the gluing so that the graph remains r-outerplanar, that is, it results in a graph embedded without edge crossings such that each vertex has a face path of length at most r to the outer face. For this we define a signature for each edge bipartition and we keep only the largest subsequence of edge bipartitions that have the same signature.

Expanding on the definition of signatures, we use it to ensure that the layer of each vertex in $G\langle A_i\rangle \circ G\langle B_j\rangle$ only decreases in comparison to G. For this, we note in the signature for each face touched by the noose that corresponds to (A_i, B_j) how far it is away from the outer face (or, rather, the face in the sphere corresponding to the outer face in the plane), and we note for each pair of faces touched by the noose how far they are away from each other. Then, if two edge bipartitions have the same signature, each vertex in the glued graph will be at most as far away from the faces touched by the noose and, hence, at most as far away from the outer face.

As we will see below, each edge-bipartition signature can be encoded in $(32r^2 + 8r) \cdot \log(r + 1) + 1$ bits. Thus, out of the $2\log(n)$ edge bipartitions that we obtain from the longest path in the decomposition tree, there are at least $\log(n)/(r + 1)^{32r^2+8r}$ edge bipartitions with the same signature.

The remainder of this section is dedicated to the formal proof of Theorem 2.

Proof (Theorem 2). In the following, fix an arbitrary r-outerplanar embedding of G.

An initial sequence $\mathcal{T}$ of edge bipartitions. Consider the canonical embedding of G into a sphere $\mathfrak{S}$ that we obtain by taking a circle that encloses but does not intersect G and identifying all points in the unbounded region of the plane which is separated off by this circle. Since G is r-outerplanar, it has branchwidth at most $2r$ [3]. By Theorem 4, there is a sphere-cut branch decomposition (T, λ) for G of width at most $2r$. We define the sequence in Theorem 2 based on (T, λ).

Consider a longest path P in T. Denote by e_1 the edge of G which is the preimage of the first vertex of P under the mapping λ. Since each edge in T induces a bipartition of the edges in G, so does each edge on P. Define the sequence $\mathcal{T} := ((C_i, D_i))_{i=1}^{t}$, where (C_i, D_i) is the bipartition of $E(G)$ induced by the ith edge on P such that $e_1 \in C_i$. We have $C_i \subsetneq C_{i+1}$ and $D_i \supsetneq D_{i+1}$

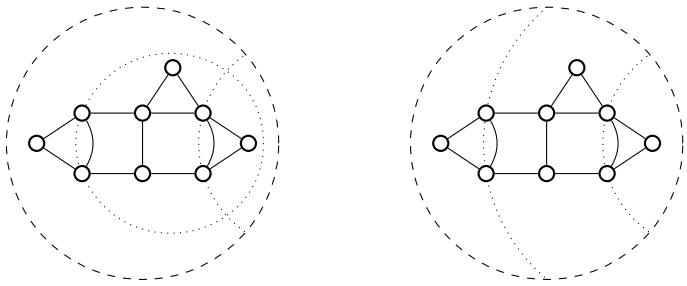

Fig. 3. A graph embedded in the sphere and two crossing nooses (dotted, left) and two noncrossing nooses (dotted, right). We projected the sphere into the plane by replacing a point in the sphere with a circle (dashed) and drawing all remaining points inside this circle. Both pairs of nooses represent the same edge bipartitions. Note that the two nooses on the right share a point on the sphere.

because T is a ternary tree and λ is a bijection. We later need a lower bound on the length of $\mathcal{T}$. For this, observe that P contains at least $2\log(n)$ edges, because G contains at least n edges (there are no vertices of degree one) and T is a ternary tree. Hence, sequence $\mathcal{T}$ also has at least $2\log(n)$ entries. The sequence in Theorem 2 is defined based on a subsequence of $\mathcal{T}$.

Obtaining a sequence of noncrossing nooses. To define the desired subsequence of $\mathcal{T}$, we choose one noose $\mathfrak{N}_i$ for each $(C_i, D_i) \in \mathcal{T}$ such that the resulting sequence of nooses has the following property. Denote by $\mathfrak{C}_i, \mathfrak{D}_i$ the open disks in which $\mathfrak{N}_i$ separates $\mathfrak{S}$ such that $C_i \subseteq \mathfrak{C}_i$ and $D_i \subseteq \mathfrak{D}_i$. Then, it shall hold that for any two i, j, $i < j$, we have $\mathfrak{C}_i \subsetneq \mathfrak{C}_j$ and $\mathfrak{D}_i \supsetneq \mathfrak{D}_j$. We say that the nooses $\mathfrak{N}_i$ and $\mathfrak{N}_j$ are *noncrossing* and *crossing* otherwise. See Figure 3 for examples.

To see that we can choose the nooses in this way, first choose them arbitrarily and then consider two crossing nooses $\mathfrak{N}_i, \mathfrak{N}_j$, $i < j$, that is, $\mathfrak{C}_i \cap \mathfrak{D}_j \neq \emptyset$. We define a noose $\tilde{\mathfrak{N}}_i$ which we obtain from $\mathfrak{N}_i$ by replacing each maximal subsegment contained in $\mathfrak{D}_j$ by the corresponding subsegment of $\mathfrak{N}_j$ which is contained in $\mathfrak{C}_i$. There is no edge of G contained in $\mathfrak{C}_i \cap \mathfrak{D}_j$ because such an edge then would also be in $C_i \cap D_j \subseteq C_i \cap D_i$, a contradiction to the fact that C_i, D_i is a bipartition of $E(G)$. Hence, noose $\tilde{\mathfrak{N}}_i$ separates $\mathfrak{S}$ into two open disks $\tilde{\mathfrak{C}}_i, \tilde{\mathfrak{D}}_i$ such that $C_i = \tilde{\mathfrak{C}}_i \cap E(G)$ and $D_i = \tilde{\mathfrak{C}}_i \cap E(G)$. Thus, $\tilde{\mathfrak{N}}_i$ fulfills the conditions for the nooses in sphere-cut branch decompositions and we may choose $\tilde{\mathfrak{N}}_i$ for (C_i, D_i) instead of $\mathfrak{N}_i$.

Clearly, $\tilde{\mathfrak{N}}_i$ and $\mathfrak{N}_j$ are noncrossing. Moreover, any noose $\mathfrak{N}_k$, $k > i$, that crosses $\tilde{\mathfrak{N}}_i$ also crosses $\mathfrak{N}_i$ because $\tilde{\mathfrak{C}}_i \subseteq \mathfrak{C}_i$. Thus, by replacing $\mathfrak{N}_i$ with $\tilde{\mathfrak{N}}_i$, the number of pairs of crossing nooses with indices at least i is strictly decreased. This means that after a finite number of such replacements we reach a sequence of pairwise noncrossing nooses.

Signatures that allow gluing. Based on the sequence $\mathcal{T}$ of edge bipartitions of G and the nooses we have fixed above for each edge bipartition, we now define a tuple, the signature, for each edge bipartition that can be encoded using $(32r^2 + 8r) \cdot \log(r + 1) + 1$ bits and that has the property that, if two edge

bipartitions have the same signature, then the corresponding graphs can be glued in a way that results in an r-outerplanar graph, as stated in Theorem 2.

We need some notation and definitions. Denote by $\mathfrak{F}$ the face in the sphere embedding of G that corresponds to the outer face of G in the planar embedding. Pick a point $y \in \mathfrak{F}$ in such a way that y is not equal to any vertex and not contained in any edge or noose $\mathfrak{N}_i$. For every noose $\mathfrak{N}_i$ we define a bijection $\beta_i \colon M(C_i, D_i) \to \{1, \ldots, |M(C_i, D_i)|\}$ corresponding to the order in which the vertices in $M(C_i, D_i)$ appear in a traversal of $\mathfrak{N}_i$ that starts in an arbitrary point. We furthermore define a map γ_i from each face touched by $\mathfrak{N}_i$ to its occurrences in the traversal of $\mathfrak{N}_i$ above. More precisely, if face $\mathfrak{G}$ occurs in the traversal of $\mathfrak{N}_i$ between vertex $\beta_i^{-1}(j)$ and $\beta_i^{-1}(j+1)$ (wherein we set $|M(C_i, D_i)| + 1$ equal to 1), then $j \in \gamma_i(\mathfrak{G})$. Finally, say that a face path P is *contained* in a closed disk $\mathfrak{E}$ if each vertex in P is contained in $\mathfrak{E}$.

The *signature* of (C_i, D_i) is a tuple (b, L_1, L_2) defined as follows.

- $b = 1$ if $y \in \mathfrak{C}_i$ and $b = 0$ otherwise.
- L_1 is the set containing the tuple $(k, \xi, \mathfrak{X}, \ell)$, for each $k \in \{1, \ldots, |M(C_i, D_i)|\}$, for each $\xi \in \{\beta, \gamma\}$, and for each $\mathfrak{X} \in \{\mathfrak{C}, \mathfrak{D}\}$, where ℓ is the length of a shortest face path from $\xi_i^{-1}(k)$ to $\mathfrak{F}$ that is contained in $\mathfrak{X}_i \cup \mathfrak{N}_i$.
- L_2 is the set containing $(k_1, k_2, \xi, \psi, \mathfrak{X}, \ell)$, for each $k_1, k_2 \in \{1, \ldots, |M(C_i, D_i)|\}$, for each pair $\xi, \psi \in \{\beta, \gamma\}$, and for each $\mathfrak{X} \in \{\mathfrak{C}, \mathfrak{D}\}$, where ℓ is the length of a shortest face path from $\xi_i^{-1}(k_1)$ to $\psi_i^{-1}(k_2)$ that is contained in $\mathfrak{X}_i \cup \mathfrak{N}_i$.

If the paths above do not exist, or the lengths are larger than r, then put ∞ instead of the length ℓ.

Definition of the desired edge bipartition sequence. Take

$$\mathcal{S} \coloneqq ((C_i, D_i, \beta_i))_{i=1}^{s}$$

where, in a slight abuse of notation, $((C_i, D_i))_{i=1}^{s}$ is the longest subsequence of $\mathcal{T}$ in which all edge bipartitions (C_i, D_i) have the same signature. Two edge bipartitions (defined via nooses) which have the same signature are shown in Figure 3 and in Figure 4. We claim that $\mathcal{S}$ fulfills the conditions of Theorem 2.

Length of the sequence. To see that the length s of $\mathcal{S}$ is large enough, recall that sequence $\mathcal{T}$ contains at least $2\log(n)$ entries. The longest subsequence of $\mathcal{T}$ with pairwise equal signatures has length at least $2\log(n)$ divided by the number of different signatures (b, L_1, L_2). It is not hard to see that there are at most two possibilities for b, at most $(r+1)^{2r \cdot 2 \cdot 2} = (r+1)^{8r}$ possibilities for L_1, and at most $(r+1)^{2r \cdot 2r \cdot 2 \cdot 2 \cdot 2} = (r+1)^{32r^2}$ possibilities for L_2, giving an overall upper bound on the number of different signatures of

$$2 \cdot (r+1)^{8r} \cdot (r+1)^{32r^2} \;=\; 2 \cdot (r+1)^{32r^2+8r}.$$

Thus $\mathcal{S}$ has length at least $\log(n)/(r+1)^{32r^2+8r}$.

Outerplanarity number of the glued graphs. For each (C_i, D_i), $(C_j, D_j) \in \mathcal{S}$, $i < j$, we have $C_i \subsetneq C_j$ and $D_i \supsetneq D_j$. Thus to prove Theorem 2 it remains to show that $G_{ij} \coloneqq G\langle C_i \rangle \circ G\langle D_j \rangle$ is r-outerplanar. To see this, we first describe

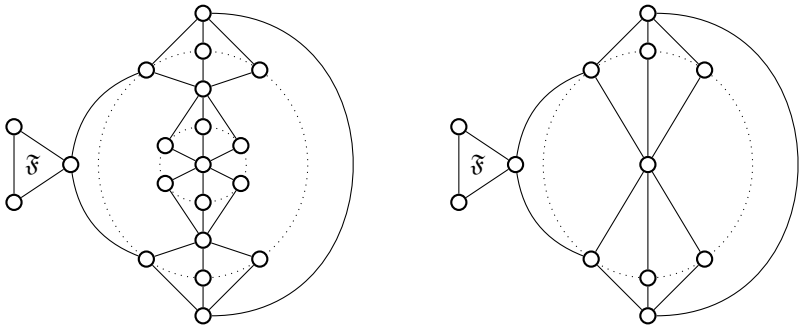

Fig. 4. Left: A graph embedded in a subdisk of the sphere which has been projected onto the plane. We show two nooses (dotted) that induce edge bipartitions. The signatures of the two edge bipartitions are the same if we assume that both left sides (the C_i) of the bipartitions contain the outermost edges in the drawing and if we furthermore assume that the corresponding mappings β_i are the clockwise orderings of the vertices on the noose starting with the topmost vertex. Right: The graph resulting from gluing along the two nooses.

how to obtain an r-outerplanar embedding for a supergraph G' of G_{ij} from G's embedding in the sphere. Graph G' is defined below and is isomorphic to G_{ij} except that it may contain multiple copies of an edge in G_{ij}.

Recall that the nooses $\mathfrak{N}_i$ and $\mathfrak{N}_j$ are noncrossing. Hence the closed disks $\mathfrak{C}_i \cup \mathfrak{N}_i$ and $\mathfrak{D}_j \cup \mathfrak{N}_j$ can intersect only in their nooses $\mathfrak{N}_i$ and $\mathfrak{N}_j$. We now consider dislocating these disks from the sphere, and identifying their boundaries $\mathfrak{N}_i$ and $\mathfrak{N}_j$, creating another sphere. Figure 4 shows an example.

Recall that the vertices in $M(C_i, D_i)$ and $M(C_j, D_j)$ are enumerated by β_i and β_j, respectively, according to traversals of the corresponding nooses. Hence, there is an open disk $\tilde{\mathfrak{C}}_i$ with $\tilde{\mathfrak{C}}_i \cap \mathfrak{C}_i = \emptyset$ and a homeomorphism $\phi\colon \mathfrak{C}_i \cup \mathfrak{N}_i \to \tilde{\mathfrak{C}}_i \cup \mathfrak{N}_j$ that has the following properties.

(i) For the two traversals of the nooses that define β_i and β_j, respectively, we have that the initial points of the traversals are mapped onto each other by ϕ and, if z comes after x in the traversal of $\mathfrak{N}_i$, then $\phi(z)$ comes after $\phi(x)$ in the traversal of $\mathfrak{N}_j$.
(ii) For each $k \in \{1, \ldots, |M(C_i, D_i)|\}$ we have $\phi(\beta_i^{-1}(k)) = \beta_j^{-1}(k)$.

Denote by G' the $\mathfrak{S}$-plane graph induced by the point set $\phi(G \cap \mathfrak{C}_i) \cup (G \cap \mathfrak{D}_i)$. We claim that from G' we can derive an r-outerplanar embedding of G_{ij}.

We first prove that G_{ij} is an edge-induced subgraph of G' without loss of generality: We may assume that G and G_{ij} have the same vertex set without loss of generality by Property (ii) of homeomorphism ϕ. Since each edge $e \in C_i$ is contained in $\mathfrak{C}_i$, it is also present in $\phi(\mathfrak{C}_i)$ and thus in G'. Moreover, each edge in $e \in D_j$ is trivially contained in $\mathfrak{D}_j$, hence, also in G'. Thus, we may assume that G_{ij} is an edge-induced subgraph of G' whence from any r-outerplanar embedding of G' we obtain an r-outerplanar embedding of G_{ij}.

Graph G' has a sphere embedding due to the way it was constructed. We now prove that from this embedding we can obtain an r-outerplanar one. This then

finishes the proof of Theorem 2. Note that there is a face in the sphere embedding of G' that contains y or $\phi(y)$ due to the flag b in the signatures. In a slight abuse of notation, we denote this face by $\mathfrak{F}$. By removing a point contained in the face $\mathfrak{F}$ from the sphere, we obtain a topological space homeomorphic to the plane. Fix a corresponding homeomorphism δ and note that, applying δ to G', we obtain a planar embedding of G' with the outer face $\delta(\mathfrak{F})$. In the following we assume that G' is embedded in this way and, for the sake of simplicity, denote $\delta(\mathfrak{F})$ by $\mathfrak{F}$.

To conclude the proof it remains to show that the embedding of G' is an r-outerplanar one. Recall that a graph is r-outerplanar if and only it has an embedding in the plane such that each vertex v has an incident face with a face path of length at most r to the outer face $\mathfrak{F}$. Call such a path *good* with respect to v.

It remains to show that each vertex in G' has a good face path. It suffices to prove this for vertices in $\mathfrak{C}_i$ whose good paths in G are not contained in $\mathfrak{C}_i$ and vertices in $\mathfrak{D}_j$ whose good paths in G are not contained in $\mathfrak{D}_j$ as the remaining ones are also present in G'. Consider a vertex in $\mathfrak{C}_i$ whose good face path P is not contained in $\mathfrak{C}_i$. We claim that we can replace every maximal face subpath of P which is contained in $\mathfrak{D}_i \cup \mathfrak{N}_i$ by a face path contained in $\mathfrak{D}_j \cup \mathfrak{N}_j$ in such a way that the resulting sequence P' is a face path in G'. Moreover, P' is at most as long as P.

Consider a maximal face subpath S of P which is contained in $\mathfrak{D}_i \cup \mathfrak{N}_i$. Each end of S is either a vertex in $M(C_i, D_i)$, or a face. If an end of S is a face, then it can either be the outer face $\mathfrak{F}$ or a face $\mathfrak{G} \neq \mathfrak{F}$ which is intersected by $\mathfrak{N}_i$. (Note that not both ends of S can be $\mathfrak{F}$ as P is a shortest path to $\mathfrak{F}$.)

If one end of S is $\mathfrak{F}$, then associate with S a tuple $(k, \xi, \mathfrak{D}, \ell)$ where $\xi = \beta$ if the other end of S is a vertex and $\xi = \gamma$ otherwise, and where ℓ is the length of S. The first entry, k, is an integer equal to $\xi_i^{-1}(v)$ if the end of S is a vertex v, and otherwise, if the end is a face $\mathfrak{G} \neq \mathfrak{F}$, then k is defined as follows. Draw an arc $\mathfrak{A}$ contained in $\mathfrak{G}$ between the two vertices that P visits before and after $\mathfrak{G}$ such that $\mathfrak{A}$ and $\mathfrak{N}_i$ have the smallest-possible intersection. Note that $\mathfrak{A}$ and $\mathfrak{N}_i$ intersect in precisely one point y since S is maximal. Define $k \in \mathbb{N}$ such that in the traversal of $\mathfrak{N}_i$ that defines β_i vertex $\beta_i^{-1}(k)$ comes before y and $\beta_i^{-1}(k+1)$ comes after y (where we set $k+1 = 1$ if $k = |M(C_i, D_i)|$).

There is a tuple $(k, \xi, \mathfrak{D}, \ell')$ with $\ell' \leq \ell$ saved in L_1 of the signature of (C_i, D_i), since S has length at most r. Thus, $(k, \xi, \mathfrak{D}, \ell')$ is also saved in L_1 of the signature of (C_j, D_j) since the signatures of (C_i, D_i) and (C_j, D_j) are the same. Hence, there is a face path S' in $\mathfrak{D}_j$ with the ends $\mathfrak{F}$ and $\xi_j^{-1}(k)$.

We claim that $\xi_j^{-1}(k)$ and $\xi_i^{-1}(k)$ describe the same entities in G'. Indeed, if $\xi = \beta$, that is, the end of S is a vertex, then $\xi_i^{-1}(k) = \beta_i^{-1}(k)$ which is equal to $\beta_j^{-1}(k) = \xi_j^{-1}(k)$ by Property (ii) of homeomorphism ϕ.

If $\xi = \gamma$, then consider the face $\mathfrak{G} = \xi_i^{-1}(k)$ and the face $\mathfrak{H} = \xi_j^{-1}(k)$, both in G. By definition, $\mathfrak{G}$ intersects $\mathfrak{N}_i$ in the segment $\mathfrak{S}_i$ of the traversal defining β between $\beta_i^{-1}(k)$ and $\beta_i^{-1}(k+1)$. Similarly, $\mathfrak{H}$ intersects $\mathfrak{N}_i$ in the segment $\mathfrak{S}_j$ between $\beta_j^{-1}(k)$ and $\beta_j^{-1}(k+1)$. In G', face $\mathfrak{G}$ is represented by $\phi(\mathfrak{G} \cap (\mathfrak{C}_i \cup \mathfrak{N}_i))$ and face $\mathfrak{H}$ is represented by $\mathfrak{H} \cap (\mathfrak{D}_j \cup \mathfrak{N}_j) = \mathfrak{H} \cap (\mathfrak{D}_j \cup \mathfrak{N}_i)$. Moreover, segments $\mathfrak{S}_i$ and $\mathfrak{S}_j$ are identified by homeomorphism ϕ because of its Property (i). Hence, $\phi(\mathfrak{G} \cap (\mathfrak{C}_i \cup \mathfrak{N}_i))$ and $\mathfrak{H} \cap (\mathfrak{D}_j \cup \mathfrak{N}_i)$ are merged into one face in G'. Thus, indeed

$\xi_j^{-1}(k)$ and $\xi_i^{-1}(k)$ describe the same entities in G'. This implies that we can replace S by S' in P and the predecessors and successors of the ends of S' in P are incident with one another.

The proof that we can replace S by a corresponding path S' in P in the case that S does not have $\mathfrak{F}$ as an end is analogous to the above and omitted. Hence, replacing all maximal face subpaths of P that are not contained in $\mathfrak{C}_i$, we obtain a good path in G'. Finally, the case that the good path of a vertex in $\mathfrak{D}_j$ is not contained in $\mathfrak{C}_i$ is symmetric to the above and also omitted.

Summarizing, since each vertex in G has a good path, so has each vertex in G', meaning that G' is r-outerplanar. Since G_{ij} is an edge-induced subgraph of G', also G_{ij} is r-outerplanar. This concludes the proof of Theorem 2. □

E Omitted content from Section 4

Proof (Lemma 1). Denote the signature of (A_i, B_i, β_i) by $(\mathcal{T}_i, \phi_i, K_i)$. There are at most $2^m - 1$ twin classes in $\mathcal{T}_i$. Furthermore, for every i, j, $i < j$, we have $A_i \subsetneq A_j$, which implies $\mathcal{T}_i \subseteq \mathcal{T}_j$. Thus, either $\mathcal{T}_i = \mathcal{T}_{i+1}$ or $\mathcal{T}_{i+1}$ comprises at least one additional twin class. Since the number of twin classes can increase at most $2^m - 2$ times, the number of different $\mathcal{T}_i$ is less than 2^m. Next, there are at most 2^m choices for a twin class for each $\beta^{-1}(i) \in M(A_i, B_i)$, leading to at most $2^{m\ell}$ different possibilities where $\ell = |M(A_i, B_i)|$. For the last part of the signature, K_i, for each γ_e there are $2^{(\ell^2-\ell)/2}$ possibilities, leading to $2^{m(\ell^2-\ell)/2}$ possibilities for K_i. Since the size ℓ of the middle sets in Theorem 2 is at most $2r$, the number of possible signatures is at most $2^m \cdot 2^{2mr} \cdot 2^{m\cdot(2r^2-r)} = 2^{m\cdot(2r^2+r+1)}$. □